

## RING BUFFERED NETWORK BUS DATA MANAGEMENT SYSTEM

This invention was made with Government support under contracts NAS4-50071 and NAS4-97010 awarded by the National Aeronautics and Space Agency. The Government has certain rights in the invention.

### FIELD OF THE INVENTION

The present invention pertains to hardware and software network systems for managing and distributing high-speed, real-time and archived data between diverse local and remote data sources, data sinks and observers.

### BACKGROUND OF THE INVENTION

The tremendous growth in networked-based distributed computing environments has created correspondingly growing needs for efficient mechanisms by which to share information. In the discussion which follows, the terms data "source" and "sink" refer to the supplier and destination of information, respectively.

Networked communication technology has historically been point-to-point. That is, a source of information and a destination (sink) for that information send and receive data over a dedicated connection between them. More recent technologies have expanded upon this concept to allow various "multicast" capabilities, wherein a data source can transmit information to a non-specific destination, and multiple data sinks can simultaneously subscribe to this flow of information through connection to a source-specific multicast address.

Additional prior art development has further abstracted the multicast concept to a "publish/subscribe" approach which allows a data source to publish data which is encoded by "subject," such that data sinks can subscribe to information by data type as opposed to a specific data source. One such technology identified by the trademark TIB, is provided by Tibco, Inc. of Palo Alto, Calif., and is described in U.S. Pat. No. 5,557,798 to Skeen et al. ("the '798 patent"). The technology described in the '798 patent features a message distribution system that uses middleware to provide the ability for data sources (publishers) to send data, and data sinks (subscribers) to request data by subject type. The '798 patent does not provide a direct means to access data from an intermediate buffered server by a combination of time-stamp and explicit source (i.e., channel) reference.

Whereas multicast and publish-subscribe approaches provide a single-source to multiple-sink (one-to-many) capability, the inverse is not generally true. That is, a many-to-one system has additional complexity and requirements. Finally, the more general many-to-many system can require an exponential growth in the number of network connections, resulting in inefficient communication performance. This many-to-many network data distribution problem has many proposed prior art solutions, each with varying degrees of efficiency and practicality.

One approach for data distribution in a many-to-many connected network is described in U.S. Pat. No. 5,634,010 to Ciscen et al. ("the '010 patent"). The '010 patent discusses intermediate data routers which can send data to multiple destinations, per predetermined routing requests. The data objects that are sent each have a time stamp, properties, addressing information, and underlying data. Application processes register for certain "types" of data, and routers send out data based on registered "interest tables." As with other related art, the '010 patent does not address time-stamped buffering of data, only routing it "on the fly."

Existing prior art addresses time-stamped data to varying degrees. Many ignore any time-relevancy of the data. Others recognize and provide means to supply "real-time" information. In both of these cases, it is left to the data sinks to receive data as it is sent from a data source; that is, the time-stamp of the data is either "as provided" or "now." In both cases however, data sinks have no efficient means to request and receive data with intermediate time-stamps, i.e., smoothly transitioned between current and historical data at the discretion and request of the data sink.

For example, a data sink may require (say for purposes of statistical analysis) the most recent contiguous block of time-series data from a source. With existing prior art technologies, the sink either has to have already saved the historical information, or ask the data source to re-transmit the data. Neither of these options are practical or convenient for a sink which lacks the required foresight or resources.

U.S. Pat. No. 4,884,217 to Skeirik describes an expert system having three classes of rules and makes brief reference to a historical process database" (HPD) with the stated goals of offering buffered, time-stamped storage to multiple sources and end-users with a "standardized" data interface. But the HPD is a small piece in a more complex expert system, and as such is not an independent network-accessible data server. Also, reference to function and structure of the HPD is described in such cursory form that it is believed one skilled in the art would be unable to use the HPD concept as a component of general data distribution and management system.

U.S. Pat. No. 5,606,705 to Randall et al. ("the '705 patent") discusses a current value table approach in which "data items" are retained until they become "too old," then they are re-fetched from the source(s). Here, data sources are queried for data, which limits the flexibility of data communication insofar as data sources are apparently not able to send data to data sinks at their own initiate. In addition, the system of the '705 patent is apparently not geared for continuous data. Instead, the system focuses on groups of quasi-static data items which may change only infrequently. As such, it appears the system of the '705 patent is geared more towards hardware or a microcontroller solution versus a portable software solution.

### SUMMARY OF THE INVENTION

One aspect of the present invention is a data management system comprising at least one server having cache memory and archive memory, and at least one RBNB module operatively associated with the at least one server. Each of the RBNB modules includes a plurality of ring buffer objects. Each ring buffer object has a ring buffer data structure for storing data, a single data input connectable to a data source and connected to the ring buffer data structure so that data provided by the data source may be stored in the ring buffer data structure, and one or more RBO data outputs connected to the ring buffer data structure for providing data stored in the ring buffer data structure. Each RBNB module also includes a plurality of network bus objects. Each network bus object has one or more inputs, each connected with at least one of the one or more RBO data outputs so as to receive data stored in the ring buffer data structure, and a single NBO data output, connectable to a data sink, for providing to the data sink the data received from the at least one of the one or more RBO data outputs.

Another aspect of the present invention is a server network for use in communicating data from at least one data source to at least one data sink, each data source having at